Capacity—or Capacitor?

Because TTL circuits are noisy during switching, most designers hang capacitors across the power busses. This General Automation design takes a different tack to supply current *capacity* without *capacitance*.

NEAL YOUNG, General Automation, Inc.

With the boon of inexpensive high-speed logic comes the bane of larger circuit boards and their associated structural, power distribution and signal interconnecting problems. Solving the power distribution problem is a technique that reduces component count and cost while increasing the effectiveness of high speed logic and, at the same time, improving structural packaging.

TTL/MSI logic meets a long-standing need: that of performing more logic operations, quicker, in a smaller space. Yet, the solution produces secondary problems to plague the designer.

In order to use TTL speed effectively, signal lead lengths must be short, minimizing propagation time. This leads to large circuit boards containing 160 to 200 ICs plus many other circuit components. Since power and ground must be distributed to all ICs, short signal paths are difficult to achieve. In addition, the larger circuit boards generally need structural support to prevent excessive warp.

Further, the TTL configuration inherently causes a high transient current condition during switching. This pulls down the V_{cc} line and pulls up the ground buss. These transients effectively reduce the noise threshold of associated circuits. Well, noise is noise, and most designers fall back on the age-old technique of suppressing noise by hanging capacitors across the power busses.

In short, the familiar solutions have resulted in exchanging one economic headache for another. Structural support of large circuit boards results in increased fabrication and production costs while noise suppression, through increasing component count, has jeopardized reliability. The cost/performance tradeoffs are very much in doubt.

The dilemma has been resolved by a patented idea in which power bussing is accomplished by means of elevated buss bars that lie in planes perpendicular to the plane of the circuit board. The power buss bars provide structural support for the circuit board while supplying power and ground to each circuit and avoiding all interference with signal lead distribution.

Used in General Automation's SPC-12, SPC-16 and System 18/30 computers, the buss bars consist of flat copper strips with connector tabs extending into the circuit board. When soldered, the tabs serve as both mechanical and electrical connections to the board. The buss bars are used in either a single layer or several overlapping layers, each having different elevations. Mechanically, the elevated buss bars provide support for boards as large as 13 by 20 inches, allowing closely spaced MSI circuits to take advantage of the fast switching time of TTL.

The conventional practice of suppressing TTL noise with capacitors also is minimized by the elevated buss bars. Because they are buss bars rather than PC runs, they supply

(Continued)

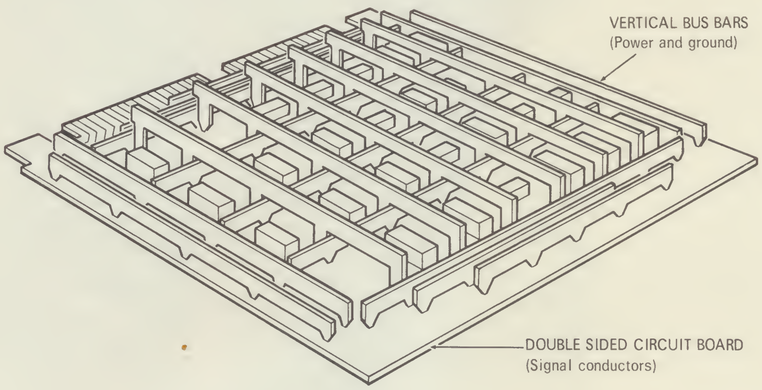

Elevated buss bars consist of parallel metal strips in multiple levels connected to a conventional two-sided circuit board. Buss bars provide both structural support for the large circuit boards and superior, lower-impedance power distribution to each MSI circuit. In addition to allowing greater area and freedom for signal lines, the elevated buss bars provide a grid shield that effectively attenuates radiated noise.

(Continued)

Capacitor? (Cont'd)

"solid" V_{cc} and ground potential at each MSI circuit, thereby curbing supply line transients (and noise) during switching. The technique reduces the need for noise suppression at each individual circuit, and it also forms a grid shield which further reduces radiated noise.

The elevated buss bar has been patented by Lawrence A. Goshorn, President of General Automation, Santa Ana, Calif. ▭

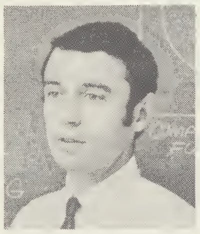

Neal Young, Manager, Manufacturing Engrg., has been responsible for the design of the SPC-12 and SPC-16 computers manufactured by General Automation. He holds a B.S.E.E. from the University of Arizona.

Why Not Capacitors?

Face it, capacitors are an important component – but they should be carefully considered when used for noise suppression. The "real-world" fact that all additional components increase the cost of the system and reduce reliability is well known and often considered as a function of design. Worst-case analysis also is well established, but the problems of field maintenance increasingly are being factored into the considerations for overall design.

Capacitors, if damaged when circuit boards are inserted or removed, can become open or shorted. Taken in its worst case, a shorted capacitor can require the opening of as many as 150 to 250 leads to find the problem. With high-speed logic, an open capacitor might result in pattern-sensitive noise. It comes and goes, and its effect is hard to trace or pin down. This can result in inestimable troubleshooting hours.

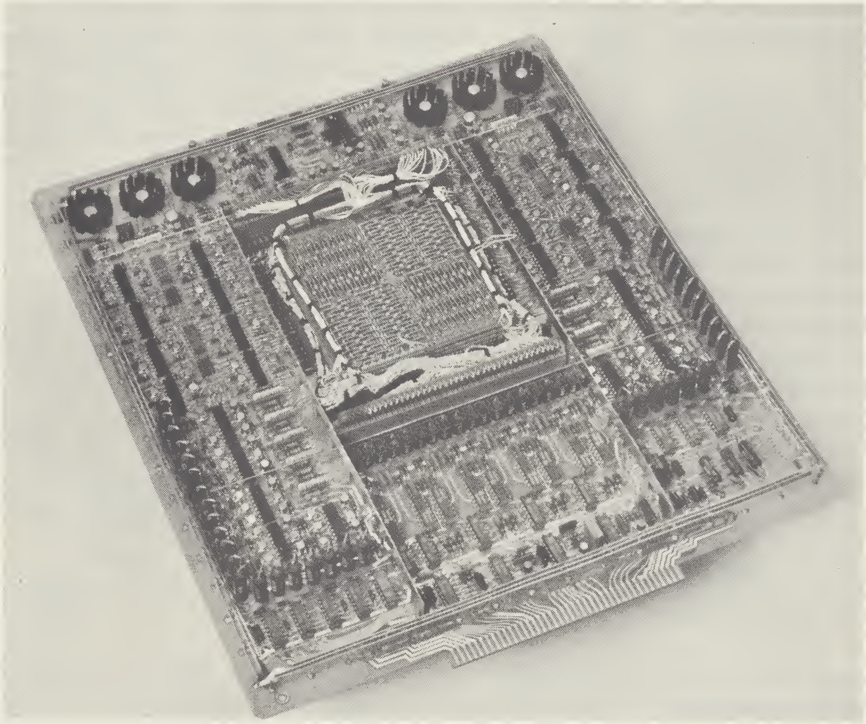

Double-layer buss bars provide power to MSI circuits on a General Automation SPC-12 memory board.

REPRINTED FROM
EDN
EXCLUSIVELY FOR DESIGNERS
A CAHNERS PUBLICATION

POWER AND PACKAGING ARE KEYS TO SEMICONDUCTOR TAKEOVER

"In dedicated computers (minicomputers) the days of cost-per-bit comparisons of memory are coming to an end," says Lawrence A. Goshorn, president of General Automation. As core sizes and prices come down to compete with semiconductor memories, the inherent power requirements of the core memory will drive the effective cost per bit back up.

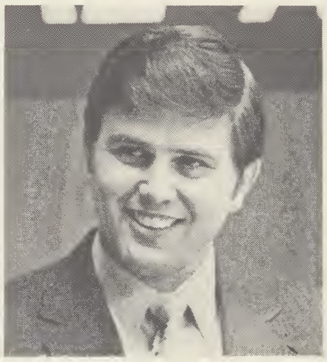

Decisions to shift from cores to semiconductor memories in small machines will be dictated by power and packaging constraints. Cost per bit of core is less and less a valid argument for its use.

Up to this point, we have been able to ignore the ferrite core's inherent physical limitations and concentrate on the economies of production. But today the small control computer and minicomputer force new thinking when it comes to using cores in this limited space. With breakthroughs in core diameter (allowed by newer memory organizations) and increasing speed of cores, keep in mind this maintains or increases the amount of power required. Of course, this means distributing power, heat, current and packaging to facilitate these equal to or greater power requirements. Raw power supply costs are tied to raw material cost such as iron; therefore, the cost per watt is difficult to decrease even in today's technology.

I see a point in the near future when the power required to drive core memory will be inconsistent with the size and packaging constraints of minicomputers. When that happens, the total system design will call for a semiconductor memory. And, at this point in time, the core memory people are going to find that from a systems viewpoint the use of semiconductor memories is cost justified. Part of our trouble is that we have grown so accustomed to cost per bit of cores going down and down that we have difficulty seeing when that cost per bit begins to go up again (from total systems cost). It does though just the minute you begin to spend money for watts and heat dissipation in tight quarters like those of the ubiquitous minicomputer.

What about the big machines, then? Although the foregoing power/package constraints won't be a factor, cost per answer will. Increasingly, cost per answer is the dominant factor in configuring these large computers. Accordingly, semiconductor memories in comparable speed will propel memories made from them into large computer use first. This push started in 1970, when giants like IBM and RCA announced semiconductor memory large computers.

P.S. A comment for the semiconductor memory people—techniques for minimizing the standby power per bit in semiconductor memories will be a necessity for use of large numbers of bits in small dedicated computers. This constraint will be primarily to the cost of standby power while the power system is shut down.

REPRINTED FROM

EDN
EXCLUSIVELY FOR DESIGNERS

A CAHNERS PUBLICATION